arXiv:1406.2545v1 [cs.NE] 10 Jun 2014

A Flexible Fitness Function for Community Detection in Complex Networks

Fabrício Olivetti de França[1] and Guilherme Palermo Coelho[2]

[1] CMCC, Federal University of ABC (UFABC) – Brazil
folivetti@ufabc.edu.br
[2] School of Technology (FT), University of Campinas (Unicamp) – Brazil
guilherme@ft.unicamp.br

Abstract. Most community detection algorithms from the literature work as optimization tools that minimize a given *fitness function*, while assuming that each node belongs to a single community. Since there is no hard concept of what a community is, most proposed fitness functions focus on a particular definition. As such, these functions do not always lead to partitions that correspond to those observed in practice. This paper proposes a new flexible fitness function that allows the identification of communities with distinct characteristics. Such flexibility was evaluated through the adoption of an immune-inspired optimization algorithm, named cob-aiNet[C], to identify both disjoint and overlapping communities in a set of benchmark networks. The results have shown that the obtained partitions are much closer to the ground-truth than those obtained by the optimization of the modularity function.

1 Introduction

Complex problems from a wide range of fields can be modeled and described as complex networks [1]. In such networks, nodes that present similar properties often tend to be linked to each other, thus forming consistent subgraphs with dense interconnections that are called *communities* [1]. The detection of communities in complex networks is an important step in the multitude of possible analyses that can be performed to such models. When a complex problem is modeled as a network, the identification of communities may allow both the comprehension of characteristics that are specific to subgroups of nodes and of how such nodes interact to each other [1].

Although communities play an important role in the network analysis, as they allow the identification of functional properties of a group of nodes and also of the ways complex behaviors emerge from simple individual functions, the formal definition of communities is still vague in the literature. Therefore, one of the challenges associated with the development of community detection algorithms is to select which metric should be used to properly evaluate whether a given set of nodes actually represent a community with characteristics that are relevant to the context of the problem. Such metrics become even more important if it is considered that a large part of algorithms for community detection are based on the optimization of *quality functions* (or *fitness* functions).

From all quality metrics described in the literature, one of the most adopted is Modularity [1], which assumes that a community is a *module of the network* and that two nodes belonging to the same community tend to have much higher probability of being connected to each other than that of two nodes belonging to different communities.

Given that each quality function may be intended to identify a set of communities according to one of the different existing definitions, the resulting partition of the original complex network invariably reflects the characteristics of such definitions. Besides, it is known that the optimization of some of these quality functions may lead to partitions of the network that do not correspond to the real partition observed in practice [2].

The above scenario is even worse when *overlapping communities* are considered. Differently from disjoint communities of complex networks, in which each node belongs to a single community, the partition of the network into overlapping communities allows some nodes (known as *bridges*) to belong to different communities at the same time. In this context, aspects such as the *clustering coefficient*[3] of the network and the number of edges connecting a given node to its neighbors (both in and out of its community) may have different impacts when choosing a bridge in complex networks that model different real-world situations.

Therefore, in this paper a new flexible fitness function for community detection is proposed. This new metric, named *Flex*, allows the user to predefine which characteristics should be present in the communities that will be obtained by the optimization process, thus allowing the identification of distinct sets of communities for the same complex network by simply adjusting a few intuitive parameters. Flex was combined with an adapted version of the immune-inspired optimization algorithm named cob-aiNet[C] [3] and applied to identify both disjoint and overlapping communities in a set of eight artificial and four real-world complex networks. The obtained results have shown that the partitions obtained with the optimization of this new metric are more coherent with the known real partitions than those obtained with the optimization of modularity.

This paper is organized as follows. The new flexible objective function for community detection will be presented in Section 2, together with some insights about how it can be applied to identify overlapping communities and a brief description of the optimization algorithm adopted here. The experimental methodology and the obtained results will be discussed in Section 3. Finally, concluding remarks and indications for future work will be given in Section 4.

2 A Flexible Objective Function for Community Detection

In a broader definition, a community structure of a network is a partition of the nodes so that each partition is densely connected. The modularity metric [1] tries to capture this definition by analyzing the difference between the number of edges

[3] Clustering coefficient corresponds to the ratio between the number of triangles formed by a given node and all possible triangles that could be formed.

inside a community and the expected number of edges that would be observed if this community was formed in a random network. Although modularity is widely adopted by the complex network community, its structure may lead to the false assumption that the number of edges between two groups decreases as the network size increases. Therefore, for larger networks, a simple connection between two nodes of different communities may result in the merging of these two communities, in order to increase (maximize) modularity. This aspect is known as the *resolution limit* of the metric [4].

Additionally, in a situation in which a given node has few links connecting it to a small community and most of its links connecting it to a large community, the optimization of modularity will often include such node into the larger community, without considering the local contribution of this node to the smaller communitiy. This can be a drawback if the node has a higher clustering coefficient with respect to the smaller community than to the larger one.

With that in mind, a new quality function for community detection, hereby called *Flex*, is proposed. The optimization of Flex tries to balance two objectives at the same time: maximize both the number of links inside a community and the local clustering coefficient of each community. Additionally, it also penalizes the occurrence of open triangles (i.e., it minimizes the *random model* effect [5]).

The first step to calculate Flex for a given partition of the network is to define the *Local Contribution* of a node i to a given community c:

$$LC(i,c) = \alpha * \triangle(i,c) + (1-\alpha) * N(i,c) - \beta * \wedge(i,c), \tag{1}$$

where $\triangle(i,c)$ is the ratio between the transitivity of node i (number of triangles that i forms) inside community c and the total transitivity of this node in the full network, $N(i,c)$ is the ratio between the number of neighbors node i has inside community c and its total number of neighbors, and $\wedge(i,c)$ is the ratio between the number of open triangles in community c that contain node i and the total participation of i in the whole network. Variables α and β are weights that balance the importance of each term.

The weight parameter α directly dictates whether the optimization process will tend to insert a given node into a clustered community or into a community that contains the majority of this node's neighbors. It is important to consider both transitivity and the neighborhood of each node in the optimization, as both concepts are not necessarily related (i.e. a given node will not necessarily have high transitivity with the majority of its neighbors). Therefore, by weighting these two criteria, the user can emphasize each of them according to what is desirable in a given practical situation.

Given the Local Contribution of all nodes to each community, the *Community Contribution* (CC) of a community c in a given partition is defined as:

$$CC(c) = \sum_{i \in c} LC(i,c) - \frac{|c|}{|V|}^{\gamma}, \tag{2}$$

where $|.|$ is the number of elements in a set, V is the set of nodes of a network and γ is the penalization weight. The penalization in this equation is devised to

avoid the generation of a trivial solution, in which the entire network forms a single community. Finally, the Flex value of a given partition p is given by:

$$Flex(p) = \frac{1}{|V|} \sum_{c \in p} CC(c) \tag{3}$$

It is also important to highlight that the penalization term of Eq. 1 ensures that the convergence to the random model is penalized, even if α favors only the number of neighbors (i.e., $\alpha = 0$) in the definition of the communities.

2.1 Applying Flex to Identify Overlapping Nodes

The Flex fitness metric also provides insights about overlapping nodes. As the importance of transitivity and neighborhood is balanced by α, this can be exploited to infer whether a given node also belongs to another community.

When using Flex as a fitness function, some nodes may be more sensible than others regarding the weight α. This happens when, for example, a node has a fraction of its neighbors on a clustered community and the remaining neighbors spread across one or more communities with lower transitivity.

Therefore, a simple heuristic that allows the identification of overlapping nodes is to search for nodes that do not make significant contribution to one of the α-weighted factors, i.e., nodes that are sensible to changes of α. After finding those nodes, we can allocate them to other communities that share a certain fraction of neighbors with them. This heuristic is summarized in Alg. 1.

Algorithm 1: Heuristic to find overlapping nodes.

Data: thresholds $thr\triangle$ and $thrN$ for the contribution to transitivity and neighborhood, respectively, and threshold $thrSh$ of shared neighbors between communities.

Result: New set of communities with overlapping nodes.

```
for each node i do
    c = community that contains i
    if △(i, c) < thr△ or N(i, c) < thrN then
        for ∀c_j ≠ c do
            if N(i, c_j) > thrSh then
                Add i to community c_j
```

2.2 The cob-aiNet[C] Algorithm

As previously mentioned, an adaptation of the cob-aiNet[C] algorithm (*Concentration-based Artificial Immune Network for Combinatorial Optimization* – [3]) was adopted in this paper to obtain a set of communities for complex networks that maximize the new proposed quality function (Flex).

The cob-aiNet[C] algorithm, which was originally proposed to solve combinatorial optimization problems [3], was previously adapted to identify both

disjoint and overlapping communities in complex networks [6]. As most of the adaptations proposed in [6] were adopted here as well, only a brief explanation of the general aspects of cob-aiNet[C] will be presented here, together with details about those aspects that differ from the adaptation proposed in [6]. For further details, the reader is referred to [3,6].

The cob-aiNet[C] is a bioinspired search-based optimization algorithm that contains operators inspired in the natural immune system of vertebrates. Therefore, it evolves a population of candidate solutions of the problem (cells or possible partitions of the complex network), through a sequence of cloning, mutation and selection steps, guided by the fitness of each individual solution.

Besides these evolutionary steps, all the cells in cob-aiNet[C] population are compared to each other and, whenever a given cell is more similar to a better one than a given threshold, its concentration (a real value assigned to each cell) is reduced. This concentration can also be increased according to the fitness of the cell (higher fitness leads to higher concentration). Such concentration-based mechanism is an essential feature of the algorithm, as it controls the number clones that will be generated for each cell at each iteration, the intensity of the mutation process that will be applied to each clone and when a given cell should be eliminated from the population (when its concentration becomes null).

When compared to the adaptations made in [6], the only differences are associated with the new hypermutation operator, which will be discussed next, and the new approach to obtain overlapping communities described in Sect. 2.1.

The new Hypermutation Operator To properly explain the new hypermutation operator, it is important to know that each cell in the population of the algorithm is represented as an array of integers with length equal to the number N of nodes of the complex network. Each position i of the array corresponds to a node of the network and assumes value $j \in \{1, 2, \ldots, N\}$ that indicates that nodes i and j belong to the same community.

This operator, which is applied to all cells in the population at a given iteration, is basically a random modification of the integer values in n_{mut} positions of the array (cell), being n_{mut} inversely proportional to the fitness and concentration of each cell (further details can be found in [3,6]).

The n_{mut} positions of the cell that will suffer mutation are randomly selected, and so are the values that will be inserted into these positions. However, the probability that a given value k (associated with node k) replaces the current value in position i is directly proportional to $|N(i) \cap N(k)|$, where $N(i)$ is the set of nodes that are neighbors of i.

3 Experimental Results

In order to assess whether Flex is able to lead to improvements in community detection, when compared to Modularity, an extensive experimental setup, composed of 8 artificial and 4 real world networks was devised. The artificial networks, which were generated by the toolbox provided by Lancichinetti [7], are composed by 4 networks formed by high density communities (i.e. with a high

number of internal edges), which facilitates the identification of the optimal partition, and 4 networks with noisy communities (i.e. with higher probability of presenting edges connecting them to other communities).

The artificial networks were generated with 50, 100, 200 and 500 nodes, being these nodes with average degree of 10 and maximum degree of 15. Such networks were generated with a maximum of 10 communities, 3 overlapping nodes belonging to an average of 2 communities and average clustering coefficient of 0.7. The mixing parameter, which introduces noise to the network structure, was set as 0.1 for the first set of networks (labeled Network $50 - 500$ in the tables that follow) and 0.3 for the second set (labeled Noise Network $50 - 500$).

The real world networks (with known partitions) that were chosen for the experiments were: Zachary's Karate Club [8]; Dolphins Social Network [9]; American College Football [10]; and a network of co-purchasing of books about US politics compiled by Krebs [11].

A total of 20 repetitions were performed for each network and for each fitness function adopted here (Flex and Modularity). The cob-aiNet[C] algorithm was empirically adjusted with the following parameters for all of the experiments: $\sigma_S = 0.2$, maximum number of iterations equal to $1,500$, $\alpha^{Ini} = 10$, $\alpha^{End} = 1$, initial population with 4 candidate solutions and maximum population size of 6. After each run, the heuristic presented in Alg. 1 was applied to the best solution returned by cob-aiNet[C].

The results were evaluated by the average of the Normalized Mutual Information, which indicates how close a given partition of the network is from the real partition (ground truth) [2,6]. In this work, the solutions with non-overlapping (labeled NMI in the tables that follow) and overlapping (labeled NMI OVER.) partitions were evaluated. Besides, the obtained results for overlapping partitions were also compared to those obtained with the technique proposed in [6] (labeled NMI MULTIMODAL).

In the following tables (Tables 2 to 4), it is also reported the evaluated fitness of the returned solutions (labeled FIT), the number of overlapping communities obtained with the proposed heuristic (labeled # Over.) and with the technique described in [6] (labeled # Over. Multimodal), the number of communities found (labeled # Comm.) and the total time in seconds taken (labeled TIME)[4].

All the parameters required by Flex were empirically defined here for groups of networks with similar structures, and are presented in Table 1. Notice though that, in practice, these parameters should be set depending on the goal of the network analysis (e.g. if partitions with highly clustered communities are required, α should be set with higher values).

Also, if the partition structure of the network is known *a priori*, the calibration of such parameters in order to obtain the known partition could indicate some characteristics of the network dynamics, such as, for example, the way that the connections of each node were established.

[4] All the experiments were performed on an Intel Core i5 with 2.7GHz, 8GB of RAM and OSX 10.9.2.

Table 1. Weight parameters and heuristic thresholds for each dataset.

Network	α	β	γ	$thr\triangle$	$thrN$	$thrSh$
Network 50, Karate	0.8	0.3	2	0.3	0.6	0.25
Network 100-500, Krebs	0.8	0.3	4	0.3	0.7	0.25
Noise Network 50-500	0.5	1.0	4	0.3	0.7	0.45
Football	0.8	0.6	4	0.3	0.6	0.25
Dolphins	0.4	0.3	4	0.3	0.6	0.25

Table 2. Results for the artificial networks with overlap and low level of noise.

	Flex	Modularity
FIT:	0.81	0.60
NMI:	0.94	0.94
NMI OVER.:	1.00	1.00
NMI MULTIMODAL:	0.87	**0.93**
# Comm.:	5.00	5.00
# Over.:	3.00	3.00
# Over. Multimodal:	4.50	0.95
TIME:	48.52	32.73

(a) Network 50

	Flex	Modularity
FIT:	0.81	0.59
NMI:	0.95	0.95
NMI OVER.:	0.98	0.98
NMI MULTIMODAL:	**0.96**	0.95
# Comm.:	4.00	4.00
# Over.:	2.10	2.00
# Over. Multimodal:	1.05	0.50
TIME:	224.33	146.19

(b) Network 100

	Flex	Modularity
FIT:	0.85	0.64
NMI:	0.97	0.97
NMI OVER.:	1.00	1.00
NMI MULTIMODAL:	**0.98**	0.97
# Comm.:	5.00	5.00
# Over.:	3.00	3.00
# Over. Multimodal:	0.65	0.00
TIME:	406.03	113.93

(c) Network 200

	Flex	Modularity
FIT:	0.87	0.80
NMI:	0.99	0.99
NMI OVER.:	1.00	1.00
NMI MULTIMODAL:	0.99	0.99
# Comm.:	11.95	12.00
# Over.:	2.95	3.00
# Over. Multimodal:	0.00	0.00
TIME:	590.99	332.31

(d) Network 500

3.1 Artificial Networks with Overlapping Communities

As expected, for the first set of networks (results given in Tab. 2), both Flex and Modularity obtained the same values of NMI for every experiment. This is due to the lower rate of noise adopted in the creation of such networks, which makes the identification of the real partitions trivial for most fitness functions. Notice though that the overlapping detection heuristic proposed here resulted in partitions with perfect NMI score (equal to 1.0) for every network considered, except Network 100. In this particular dataset, one of the overlapping nodes did not attend the criteria established by the heuristic, thus the obtained NMI was slightly lower than 1.0.

The results for the second set of networks, which were generated with higher noise, are reported in Tab. 3. In this scenario, the differences between Flex and Modularity become much more evident. In every situation the heuristic combined with Flex was able to find most of the overlapping nodes of each problem, as pointed out by the higher values of NMI. On the other hand, the method proposed in [6] obtained lower values of NMI by introducing many more (false) overlapping nodes into the partition. It is also noticeable that, in the presence of noise, Modularity also tends to find partitions with much less communities than Flex, which is due to the resolution limit discussed in Sect. 2.

It is important to highlight that the results were statistically verified by means of the Kruskal paired test with significance < 0.05. Therefore, in Tables 2–4, those results with statistical significance (p-value< 0.05) are marked in bold.

Table 3. Results for the artificial networks with overlap and high level of noise.

	Flex	Modularity
FIT:	0.47	0.55
NMI:	**0.77**	0.43
NMI OVER.:	**0.78**	0.43
NMI MULTIMODAL:	**0.69**	0.42
# Comm.:	6.00	3.00
# Over.:	3.00	0.00
# Over. Multimodal:	12.80	1.60
TIME:	58.77	73.34

(a) Noise Network 50

	Flex	Modularity
FIT:	0.51	0.49
NMI:	**0.93**	0.58
NMI OVER.:	**0.94**	0.58
NMI MULTIMODAL:	**0.82**	0.43
# Comm.:	6.00	4.30
# Over.:	0.40	1.45
# Over. Multimodal:	15.60	35.40
TIME:	143.46	200.52

(b) Noise Network 100

	Flex	Modularity
FIT:	0.51	0.50
NMI:	**0.93**	0.60
NMI OVER.:	**0.94**	0.62
NMI MULTIMODAL:	**0.81**	0.47
# Comm.:	6.00	4.00
# Over.:	2.00	5.50
# Over. Multimodal:	34.65	40.55
TIME:	655.99	657.93

(c) Noise Network 200

	Flex	Modularity
FIT:	0.50	0.63
NMI:	**0.89**	0.55
NMI OVER.:	**0.91**	0.56
NMI MULTIMODAL:	**0.73**	0.40
# Comm.:	13.95	7.65
# Over.:	10.30	6.70
# Over. Multimodal:	145.70	194.00
TIME:	1818.17	2619.40

(d) Noise Network 500

3.2 Real-world Social Networks

Regarding the real-world social networks, the obtained results (given in Tab. 4) show that, again, Flex leads to a significant improvement over Modularity. However, it is important to notice that the ground truths of such networks are related to the classification of the nodes of these datasets according to their respective domains, so it does not necessarily mean that they actually correspond to the true partitions of the networks. Therefore, it maybe practically impossible to reach a perfect NMI score in some situations. It is also important to notice that those ground truth were originally devised without overlapping, so the NMI score with overlapping will always be smaller than the original NMI score.

Some interesting characteristics of each of these networks can be identified through a combination of visual inspection of the obtained partitions together with an analysis of the weights (α and β) that led to the best values of NMI. From the obtained results, it is possible to infer that both the Karate Club and Krebs networks are formed by highly clustered communities ($\alpha = 0.8$), which makes sense as the Karate Club is a small social network prone to mutual friendships and the Krebs network, captures the interest of readers about particular subjects, so they tend to buy only books that are related to their political views.

The Football network required the same value for α but a much higher value for β, which means that this particular network does not allow open triangles. The reason for that is due to the organization of tournaments that limit the occurrence of intra-cluster relationships. Finally, for the Dolphins network the required weights are more favorable to the establishment of inter-community relationships instead of clustering, which might be related to the hierarchy in their society that favors the creation of several hubs inside a community [9], thus raising the number of open triangles.

In order to illustrate the overlapping communities obtained by the combination of Flex and the proposed heuristic, Fig. 1 depicts the best partitions with

Table 4. Results obtained for the real world networks.

	Flex	Modularity
FIT:	0.82	0.41
NMI:	**0.95**	0.40
NMI OVER.:	**0.79**	0.45
NMI MULTIMODAL:	**0.91**	0.55
# Comm.:	1.95	3.95
# Over.:	1.90	2.90
# Over. Multimodal:	0.55	16.30
TIME:	39.36	50.24

(a) Karate Club

	Flex	Modularity
FIT:	0.68	0.53
NMI:	**0.86**	0.46
NMI OVER.:	**0.82**	0.46
NMI MULTIMODAL:	**0.82**	0.46
# Comm.:	2.30	4.00
# Over.:	1.90	7.00
# Over. Multimodal:	4.60	0.00
TIME:	63.88	27.61

(b) Dolphins

	Flex	Modularity
FIT:	0.77	0.60
NMI:	**0.74**	0.67
NMI OVER.:	**0.72**	0.66
NMI MULTIMODAL:	**0.74**	0.67
# Comm.:	11.25	9.60
# Over.:	2.25	2.05
# Over. Multimodal:	0.00	0.00
TIME:	23.36	26.31

(c) Football

	Flex	Modularity
FIT:	0.72	0.53
NMI:	**0.45**	0.32
NMI OVER.:	**0.43**	0.37
NMI MULTIMODAL:	**0.45**	0.32
# Comm.:	2.25	4.95
# Over.:	7.30	11.90
# Over. Multimodal:	6.20	0.75
TIME:	126.01	47.33

(d) Krebs

overlapping nodes for each problem. In Fig. 1, the colors represent the communities found by the optimization of Flex, the shapes represent the communities in the ground truth and the larger nodes are the overlapping nodes. It is visually noticeable that the community formation found using the Flex function makes sense, given the weights for each network. Also, every overlapping node is clearly positioned between two or more distinct groups.

4 Conclusion

In this paper a novel fitness function for community detection in complex networks was introduced, together with a heuristic that, based on particular characteristics of this function, allows the identification of overlapping nodes. This new fitness function, called Flex, is parametrized in such a way that it can be adapted to obtain communities with different characteristics.

Through an extensive experimental setup, in which Flex was combined with an immune-inspired algorithm adopted for the optimization process (a novel mutation operator was also proposed in this paper), it was possible to verify that this new fitness function and heuristic are capable of leading to partitions close to the ground truth of a set of networks with different characteristics.

As for future investigations, we intend to evaluate a multi-objective version of the proposed fitness function and verify whether it leads to a better set of overlapping nodes. Besides, a thorough sensitivity analysis of Flex's parameters will also be performed.

References

1. Newman, M.E.J., Girvan, M.: Finding and evaluating community structure in networks. Proceedings of the National Academy of Sciences of the United States of America **99**(12) (2004) 7821–7826

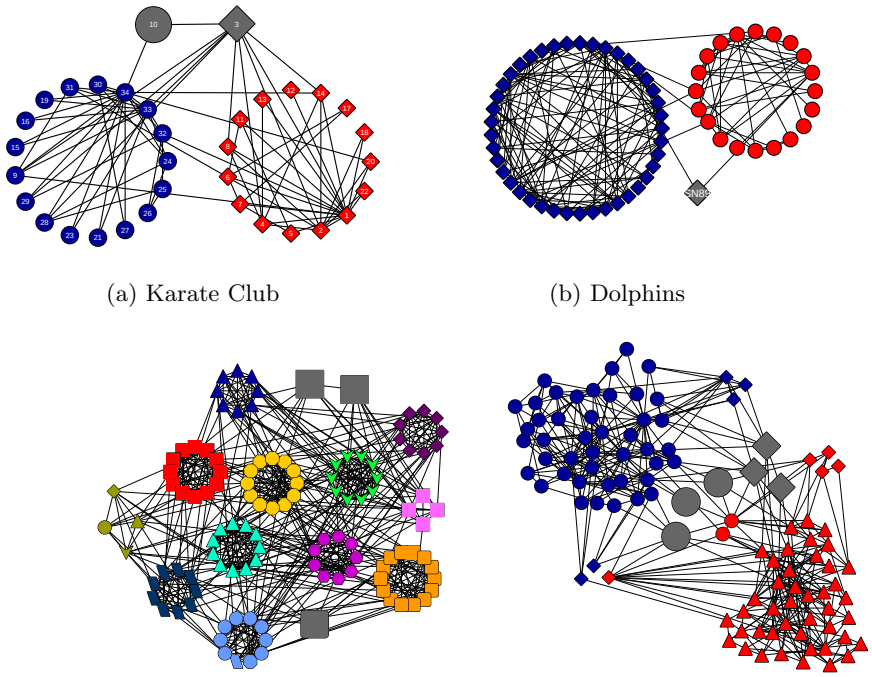

(a) Karate Club (b) Dolphins

(c) Football (d) Krebs

Fig. 1. Results obtained by cob-aiNet[C] with Flex for the real-world networks.

2. Lancichinetti, A., Radicchi, F., Ramasco, J.J., Fortunato, S.: Finding statistically significant communities in networks. PloS One **6** (2011) e18961
3. Coelho, G.P., de França, F.O., Von Zuben, F.J.: A concentration-based artificial immune network for combinatorial optimization. In: Proc. of the 2011 IEEE Congress on Evolutionary Computation (CEC). (2011)
4. Fortunato, S., Barthelemy, M.: Resolution limit in community detection. Proceedings of the National Academy of Sciences **104**(1) (2007) 36–41
5. Eerdös, P., Rényi, A.: On random graphs I. Publ. Math. Debrecen **6** (1959) 290–297
6. de França, F.O., Coelho, G.P.: Identifying overlapping communities in complex networks with multimodal optimization. In: Proc. of the 2013 IEEE Congress on Evolutionary Computation (CEC), IEEE (2013) 269–276
7. Lancichinetti, A., Fortunato, S., Radicchi, F.: Benchmark graphs for testing community detection algorithms. Physical Review E **78**(4) (2008) 046110
8. Zachary, W.W.: An information flow model for conflict and fission in small groups. Journal of Anthropological Research **33** (1977) 452–473
9. Lusseau, D., Schneider, K., Boisseau, O.J., Haase, P., Slooten, E., Dawson, S.M.: The bottlenose dolphin community of Doubtful Sound features a large proportion of long-lasting associations. Behav. Ecology and Sociobiol. **54**(4) (2003) 396–405
10. Evans, T.: Clique graphs and overlapping communities. Journal of Statistical Mechanics: Theory and Experiment **2010**(12) (2010) P12037
11. Krebs, V.: The social life of books, visualizing communities of interest via purchase patterns on the WWW (2004)